REMARKS and ARGUMENTS

Applicants appreciate the Examiners attention to this application. Claims 1-37 are pending in the Application. All claims stand rejected under 35 U.S.C. § 103(a). A declaration was previously submitted under 37 C.F.R. § 1.131 to disqualify one of the references as prior art. This response provides further specificity regarding the previously-submitted Invention Disclosure Form in relation to the claimed subject matter.

Objections to Prior-filed Declaration

With the previous response, a Supplemental Declaration was filed. Applicants maintain that the original declaration along with the Supplemental Declaration is effective to overcome the Damron reference.

Evidence of Conception. In the Advisory Action, the Examiner requests that Applicants clearly show how the Invention Disclosure Form ("IDF") proves conception by showing where each claim can be found in the document. The discussion below shows, in complete detail, that the inventors were in possession of the claimed invention prior to the effective date. Although not exhaustive regarding every place in the IDF where each claim element is disclosed, the following discussion shows in sufficient detail *at least one place* in the IDF that conception of each claim element is disclosed.

Claim 1. In particular, the section entitled "1. Executive Summary" discloses first and second processors in the first figure of that section. In the second paragraph of such section, the Applicants demonstrate that they have possession of the conception of Speculative Precomputation, which involves a "main thread". It is obvious from this section that the main thread is a main thread instruction stream – for example, the section discloses that critical data

are in the cache by the time they are "needed" (e.g., required as an input for an instruction in the main thread).

This discussion of Speculative Precomputation in the Executive Summary indicates a working knowledge of Speculative Precomputation, including a working knowledge that Speculative Precomputation involves "helper threads" that include a subset of the main thread instruction stream, where the subset includes the delinquent instruction. This implied knowledge is evidenced at the last paragraph of section 3.4. This paragraph refers to the helper thread, and how it is written to include the "delinquent loads of interest". See also Section 3, which states that "the helper threads are running prefetch slices, i.e., a subset of the main program." A "delinquent load" is understood, based on this implied knowledge of Speculative Precomputation, to be a load instruction that is predicted to miss in the cache.

The first figure in the Executive Summary discloses multiple processor cores with private caches (see second balloon from top on the left-hand side of the figure). The caches include private instruction caches (shown as "I" in the figure) and private data caches (shown as "D" in the figure).

The first figure in the Executive Summary further discloses a shared memory system – see "off-chip memory" and "on-chip Shared L3" in the figure. This shared memory is coupled to both the first and second processors (shown as shaded cores + private caches in the figure).

The second paragraph of the Executive Summary further discloses "helper-induced pre-fetches" to pre-fetch data into the cache.

The Executive Summary includes five numbered sub-paragraphs. In the first numbered sub-paragraph, it is disclosed that "when a miss from one core's cache is serviced by the shared

cache." This is evidence of conception of "logic to retrieve, responsive to a miss of requested data for the delinquent instruction in the private cache of the second processor, the requested data from the shared memory system." (See Claim 1). Such sub-paragraph further states that "the data will be injected into all cores' private data caches." This statement is evidence of conception of "the logic to provide the requested data to the private data cache of the first processor." (See Claim 1).

Claim 2. First figure of the Executive Summary (see CMP chip package).

Claim 3. The first figure in the Executive Summary "on-chip Shared L3".

Claim 4. Executive Summary – "sharing occurs higher up in the hierarchy." See also, Section 2: "but CMP models have private lower level caches and only share memory resources at higher levels of the memory hierarchy." Use of the word "levels" in the plural form indicates more than one shared cache in the memory hierarchy.

Claim 5. First figure of the Exeuctive summary, showing on-chip Shared L3 cache inside CMP chip package.

Claim 6. Section 3.1 – "when the higher level cache (such as shared L3) services a miss from one core's lower level private cache (such as private L2), it broadcasts that result data return to all helper cores."

Claim 7. First Figure of Executive Summary.

Claim 8. Section 3.1 – "when the higher level cache (such as shared L3) services a miss from one core's lower level private cache (such as private L2), it broadcasts that result data return to all helper cores."

Claim 9. Section 3.1 – return data from the cache is selectively multicast to a subset of the cores.

Claim 10. 4.3, subsection a – helper threads are spawned by main thread in response to triggers.

Claim 11. See sections above regarding first and second processors, main thread instruction stream, delinquent instruction, helper thread/subset of main thread instruction stream, first and second processors having private data caches. For the last two "logic" elements of Claim 11, please see section 3.2 of the IDF: when one core's [private data cache] and then shared cache incurs a cache miss, the "other cores' private caches opportunistically return the requested data if they have it available."

Claim 12. Inherent in defaut cache hierarchy processing.

Claim 13. "interconnect network" in third numbered subparagraph of Executive Summary.

Claim 14. Section 3.1 -"Broadcasting..."

Claim 15. First figure of Executive Summary.

Claim 16. Executive Summary – "sharing occurs higher up in the hierarchy." See also, Section 2: "but CMP models have private lower level caches and only share memory resources at higher levels of the memory hierarchy." Use of the word "levels" in the plural form indicates more than one shared cache in the memory hierarchy.

Claim 17. 4.3, subsection a – helper threads are spawned by main thread in response to triggers.

Claims 18 and 25. Executive Summary.

Claims 19 and 26. The Executive Summary includes five numbered sub-paragraphs. In the first numbered sub-paragraph, it is disclosed that "when a miss from one core's cache is serviced by the shared cache." This is evidence of conception of "retrieving the load data from a shared memory system." (See Claim 19). Such sub-paragraph further states that "the data will be injected into all cores' private data caches." This statement is evidence of conception of "providing the load data to the private cache of the main core." (See Claim 19). Evidence of conception of this latter element is also seen at the first numbered sub-paragraph of the Executive Summary, which states "helper thread prefetches to be realized much earlier in the main thread's core..."

Claims 20 and 27. First numbered sub-paragraph of Executive Summary. Such sub-paragraph further states that "the data will be injected into all cores' private data caches."

Claims 21 and 28. First numbered sub-paragraph of Executive Summary. Such sub-paragraph further states that "the data will be injected into all cores' private data caches."

Claims 22 and 29. Second numbered sub-paragraph of Executive Summary.

Claims 23 and30. Section 3.1 -"Broadcasting..."

Claims 24 and 31. providing load data to main core from shared memory system – first numbered sub-paragraph of Executive Summary: "a miss from one core's cache is serviced by the shared cache." Providing load data to main core from the private cache of one of the plurality of cores – Section 3.2.

Claim 32. Executive Summary.

Claim 33. Executive Summary – "the helper threads will help the main thread to warm up the L3 cache."

Claim 34. Section 3.2.

Claim 35. Inherent in known cache processing.

Claim 36. Third numbered sub-paragraph of Executive Summary.

Claim 37. First figure of Executive Summary.

In sum, Applicants have shown herein that the original declaration submitted by the Applicants along with the Supplemental Declaration submitted with the previous response by Applicants are sufficient to overcome the Damron reference.

Claim Rejections -35 USC § 103(a)

The Office Action has rejected Claims 1-3, 5-9, and 11-13, 15, 18-22 and 25-29 under 35 U.S.C. § 103(a) as being unpatentable over Damron (U.S. Patent Application Publication No. US 2004/0148491 A1) in view of Jamil (U.S. Patent Application Publication No. US 2003/0126365 A1). Although not specifically listed in paragraph 5 of the Office Action (the introductory paragraph for the Damron/Jamil rejections), it appears from the body of the Office Action that Claim 31 also stands rejected on these grounds (see text between paragraphs 20 and 21 of the Office Action).

The Office Action has also rejected Claims 4, 16, and 32-37 under 35 U.S.C. § 103(a) as being unpatentable over Damron in view of Jamil and in further view of Jeddeloh (U.S. Patent No. 6,789,168 B2). Finally, the Office Action also rejects Claims 10 and 17 under 35 U.S.C. § 103(a) as being unpatentable over Damron in view of Jamil and in further view of Luk (U.S. Patent Application Publication No. US 2002/0055964 A1). However, the Office Action has

failed to meet its burden of makings it prima facie case of obviousness for the claims, and such rejections should be withdrawn.

The filing date of Damron is January 28, 2003 (the "Effective Date"). The present invention was conceived before that Effective Date. The Application was diligently drafted and filed during the time between just before the filing date of Damron until July 31, 2003. Previously filed were an Original Declaration and also a Supplemental Declaration and associated evidence which provide additional facts concerning the conception and diligent constructive reduction to practice of the present invention.

In view of the foregoing, a prima facie case of obviousness has not been made with respect to any claim in the case. Since all of the rejections under 35 U.S.C. § 103(a) rely on Damron, to the extent that those rejections might be applied to the claims, those rejections should be withdrawn. All claims remaining in the case should therefore be allowed.

Accordingly, Applicants respectfully submit that the applicable rejections have been overcome and must all be withdrawn. Applicants reserve all rights with respect to the application of the doctrine equivalents. Applicant respectfully requests that a timely Notice of Allowance be issued in this case. If the Examiner feels that an interview would help to resolve any remaining issues in the case, the Examiner is invited to contact Shireen Bacon of Intel, at (512) 732-3917.

Please charge any shortages and credit any overcharges to our Deposit Account No. 02-2666.

Respectfully submitted,

Dated: December 8, 2006   /Shireen Irani Bacon/
                                 Shireen Irani Bacon
                                 Reg. No. 40,494
                                 Tel.:(512) 732-3917

                                 12400 Wilshire Boulevard
                                 Seventh Floor
                                 Los Angeles, CA 90025-1026